
Subject: Testing migration: weird behaviour of process inside VPS

Posted by [atomico80](#) on Thu, 03 Aug 2006 14:24:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

I'm testing the migration of a VPS between 2 HN.

HNa and Hnb are the real machines both running Debian and 2.6.16-026test015 kernel

then I have some VPS(with all of them the behaviour is the same) all with debian updated from the minimal template.

Now I'm running a simple counter on my VPS, just to see what happens when I migrate the server.

The counter just displays numbers on the screen.

I'm logged in all the machines with putty, even on the vps to see what kind of delay can experience the user during the migration.

That's what happens: when I migrate from HNa to HNb I almost don't see the freezing of the system and the migration is successful.

Result: from an user point of view all ok.

When I migrate a VPS(the same VPS or another, same results)from HNb to HNa, running the counter in the same situation, it happens that the counter stops as frozen without waking up, but the migration works successfully.

Checking the running processes on the destination machine(HNa now) I see the process of the counter still alive, but it is producing no number on the screen, he is blocked. To get back the control of the virtual machine i have to do ctrl+c.

did anybody else tried?

do u need some more info?

Can somebody tell me what to check and to fix to avoid that?

thanks

Ale

Subject: Re: Testing migration: weird behaviour of process inside VPS

Posted by [Andrey Mirkin](#) on Fri, 04 Aug 2006 08:02:25 GMT

[View Forum Message](#) <> [Reply to Message](#)

Can you please attach here source code of this program.

Please try to synchronize times on both machines and repeat this test.

Subject: Re: Testing migration: weird behaviour of process inside VPS

Posted by [atomico80](#) on Fri, 18 Aug 2006 09:46:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

ok the 2 HN are now sincronized.

and i get the same problem as before

the clock of the VPS has a differen time than the HN.

but the command ntpdate give me error cause It's not possible to access to the HW

Can't adjust the time of day: Operation not permitted

Anyway here is the code of the counter

```
#include <stdio.h>
```

```
int main (void) {  
    int i=1;  
    while(i<=1000000){  
        printf ("counter=%-8d\n", i);  
        sleep (1);  
        i++;  
    }  
    return 0;  
}
```

Thanks

Subject: Re: Testing migration: weird behaviour of process inside VPS

Posted by [Andrey Mirkin](#) on Fri, 18 Aug 2006 10:29:13 GMT

[View Forum Message](#) <> [Reply to Message](#)

Does the problem still exist after time is sinchronized?

If exists then please strace your program when it hangs and post here strace output.

Subject: Re: Testing migration: weird behaviour of process inside VPS

Posted by [atomico80](#) on Fri, 18 Aug 2006 11:24:01 GMT

[View Forum Message](#) <> [Reply to Message](#)

I corrected my former post.

yes same problems.

here is the strace:

```
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 8) = 0  
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 8) = 0  
rt_sigprocmask(SIG_SETMASK, [], NULL, 8) = 0  
nanosleep({1, 0}, {1, 0})           = 0
```

```

write(1, "counter=65  \n", 17counter=65
) = 17
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 8) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 8) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 8) = 0
nanosleep({1, 0}, {1, 0}) = 0
write(1, "counter=66  \n", 17counter=66
) = 17
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 8) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 8) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 8) = 0
nanosleep({1, 0}, {1, 0}) = 0
write(1, "counter=67  \n", 17counter=67
) = 17
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 8) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 8) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 8) = 0
nanosleep({1, 0}, {1, 0}) = 0
write(1, "counter=68  \n", 17counter=68
) = 17
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 8) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 8) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 8) = 0
nanosleep({1, 0}, 0xbf9c8e24) = ? ERESTART_RESTARTBLOCK (To be restarted)
setup(<unfinished ...>

```

It is of course the last part, because it's always the same, but the end...

to take the control back i have to use ctrl+c

Thanks again

Subject: Re: Testing migration: weird behaviour of process inside VPS

Posted by [whately](#) on Thu, 29 May 2008 14:17:19 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello,

I got the same error when migrating a CT running a task that use nanosleep. (the application is crashing after the migration)

The tested application is "vmstat 1"

Here is the trace:

```

execve("/usr/bin/vmstat", ["vmstat", "1"], [/* 20 vars */]) = 0
uname({sys="Linux", node="virtua0", ...}) = 0
brk(0) = 0x82e1000

```

```

access("/etc/ld.so.preload", R_OK) = -1 ENOENT (No such file or directory)
open("/etc/ld.so.cache", O_RDONLY) = 3
fstat64(3, {st_mode=S_IFREG|0644, st_size=8773, ...}) = 0
old_mmap(NULL, 8773, PROT_READ, MAP_PRIVATE, 3, 0) = 0xb7f6a000
close(3) = 0
open("/lib/libproc-3.2.3.so", O_RDONLY) = 3
read(3, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\360!\0\0004\ 0\0\0 "..., 512) = 512
fstat64(3, {st_mode=S_IFREG|0755, st_size=44736, ...}) = 0
old_mmap(NULL, 67512, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3, 0)
= 0x8a3000
old_mmap(0x8ae000, 4096, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0xa000) = 0x8ae000
old_mmap(0x8af000, 18360, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x8af000
close(3) = 0
open("/lib/tls/libc.so.6", O_RDONLY) = 3
read(3, "\177ELF\1\1\1\0\0\0\0\0\0\0\0\3\0\3\0\1\0\0\0\320N\1\0004\ 0\0\0 "..., 512) = 512
fstat64(3, {st_mode=S_IFREG|0755, st_size=1521908, ...}) = 0
old_mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1,
0) = 0xb7f69000
old_mmap(NULL, 1223900, PROT_READ|PROT_EXEC, MAP_PRIVATE|MAP_DENYWRITE, 3,
0) = 0xb7a000
old_mmap(0xc9f000, 16384, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x124000) = 0xc9f000
old_mmap(0xca3000, 7388, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0xca3000
close(3) = 0
old_mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1,
0) = 0xb7f68000
mprotect(0xc9f000, 8192, PROT_READ) = 0
mprotect(0x921000, 4096, PROT_READ) = 0
set_thread_area({entry_number:-1 -> 6, base_addr:0xb7f686c0, limit:1048575, seg_32bit:1,
contents:0, read_exec_only:0, limit_in_pages:1, seg_not_present:0, useable:1}) = 0
munmap(0xb7f6a000, 8773) = 0
uname({sys="Linux", node="virtua0", ...}) = 0
brk(0) = 0x82e1000
brk(0x8302000) = 0x8302000
open("/etc/mtab", O_RDONLY) = 3
fstat64(3, {st_mode=S_IFREG|0444, st_size=0, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7f6c000
read(3, "simfs / simfs rw 0 0\nnone /proc "..., 4096) = 120
close(3) = 0
munmap(0xb7f6c000, 4096) = 0
open("/proc/stat", O_RDONLY) = 3
fstat64(3, {st_mode=S_IFREG|0444, st_size=0, ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7f6c000

```

```

read(3, "cpu 139 0 625 754875 0 0 0 0\nncp"..., 4096) = 183
read(3, "", 4096) = 0
close(3) = 0
munmap(0xb7f6c000, 4096) = 0
ioctl(1, TIOCGWINSZ, {ws_row=43, ws_col=157, ws_xpixel=0, ws_ypixel=0}) = 0
fststat64(1, {st_mode=S_IFCHR|0620, st_rdev=makedev(136, 1), ...}) = 0
mmap2(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) =
0xb7f6c000
write(1, "procs -----memory-----"..., 79) = 79
write(1, " r b swpd free buff cach"..., 79) = 79
open("/proc/meminfo", O_RDONLY) = 3
lseek(3, 0, SEEK_SET) = 0
read(3, "MemTotal: 2072384 kB\nMemFre"..., 1023) = 771
open("/proc/stat", O_RDONLY) = 4
read(4, "cpu 139 0 625 754875 0 0 0 0\nncp"..., 8191) = 183
open("/proc/vmstat", O_RDONLY) = 5
lseek(5, 0, SEEK_SET) = 0
read(5, "nr_active 0\nnr_inactive 0\nnr_ano"..., 1023) = 721
write(1, " 0 0 0 2061928 0 "..., 81) = 81
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 0) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 0) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 0) = 0
nanosleep({1, 0}, {1, 0}) = 0
lseek(3, 0, SEEK_SET) = 0
read(3, "MemTotal: 2072384 kB\nMemFre"..., 1023) = 771
lseek(4, 0, SEEK_SET) = 0
read(4, "cpu 139 0 625 755075 0 0 0 0\nncp"..., 8191) = 183
lseek(5, 0, SEEK_SET) = 0
read(5, "nr_active 0\nnr_inactive 0\nnr_ano"..., 1023) = 721
write(1, " 0 0 0 2061916 0 "..., 81) = 81
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 0) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 0) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 0) = 0
nanosleep({1, 0}, {1, 0})
= 0
... (same lines repeated)

lseek(3, 0, SEEK_SET) = 0
read(3, "MemTotal: 2072384 kB\nMemFre"..., 1023) = 771
lseek(4, 0, SEEK_SET) = 0
read(4, "cpu 139 0 625 755476 0 0 0 0\nncp"..., 8191) = 183
lseek(5, 0, SEEK_SET) = 0
read(5, "nr_active 0\nnr_inactive 0\nnr_ano"..., 1023) = 721
write(1, " 0 0 0 2061916 0 "..., 81) = 81
rt_sigprocmask(SIG_BLOCK, [CHLD], [], 0) = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, 0) = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, 0) = 0
nanosleep({1, 0}, {1, 0}) = 0

```

```
lseek(3, 0, SEEK_SET)          = 0
read(3, "MemTotal: 2072384 kB\nMemFre"... , 1023) = 771
lseek(4, 0, SEEK_SET)          = 0
read(4, "cpu 139 0 625 757881 0 0 0 0\nncp"... , 8191) = 183
lseek(5, 0, SEEK_SET)          = 0
read(5, "nr_active 0\nnr_inactive 0\nnr_ano"... , 1023) = 721
write(1, " 0 0 0 2061916 0 "... , 81) = 81
rt_sigprocmask(SIG_BLOCK, [CHLD], [], = 0
rt_sigaction(SIGCHLD, NULL, {SIG_DFL}, = 0
rt_sigprocmask(SIG_SETMASK, [], NULL, = 0
nanosleep({1, 0},
```

it's runned on

Linux virtua0 2.6.18-53.1.13.el5.028stab053.10 #1 SMP Tue Apr 1 14:58:47 MSD 2008 i686 i686
i386 GNU/Linux

Is there any fix for this problem already ?

thanks,
Lauro.
